

Can Almost Everybody be Almost Happy? PCP for PPAD and the Inapproximability of Nash

Yakov Babichenko*

Christos Papadimitriou[†]Aviad Rubinfeld[‡]

Abstract

We conjecture that PPAD has a PCP-like complete problem, seeking a near equilibrium in which all but very few players have very little incentive to deviate. We show that, if one assumes that this problem requires exponential time, several open problems in this area are settled. The most important implication, proved via a “birthday repetition” reduction, is that the $n^{O(\log n)}$ approximation scheme of Lipton et al. [23] for the Nash equilibrium of two-player games *is essentially optimum*. Two other open problems in the area are resolved once one assumes this conjecture, establishing that certain approximate equilibria are PPAD-complete: Finding a *relative* approximation of two-player Nash equilibria (without the well-supported restriction of [14]), and an approximate competitive equilibrium with equal incomes [10] with small clearing error and near-optimal Gini coefficient.

*Technion. Email: yakovbab@tx.technion.ac.il

[†]University of California, Berkeley. Email: christos@cs.berkeley.edu

[‡]University of California, Berkeley. Email: aviad@cs.berkeley.edu

1 Introduction

It is known that finding an ε -approximate Nash equilibrium in a two-person game:

1. is PPAD-complete if ε is inversely polynomial [12];
2. but can be solved in quasipolynomial time for any fixed $\varepsilon > 0$ [23]; and
3. the smallest known polynomially attainable approximation ratio is still over $\frac{3}{10}$ [28].

These three facts articulate rather dramatically the mystery, by now almost a decade old, of the problem's approximability.

Can the inversely polynomial inapproximability bound of [12] be improved to constant? Unlikely, because by [23] this would imply a quasipolynomial algorithm for the iconic problem in PPAD:

Definition 1. END OF THE LINE: ([15]) Let S and P be two circuits with (at most) $\tilde{O}(n)$ gates each (computing the predecessor and successor correspondences) with n input bits and n output bits each, such that $P(0^n) = 0^n \neq S(0^n)$, find an input $x \in \{0, 1\}^n$ such that $P(S(x)) \neq x$ or $S(P(x)) \neq 0^n$.

Conventional wisdom, supported by black-box lower bounds [21, 5], is that this problem requires exponential time to solve. In direct analogy with the exponential time hypothesis [22], we state the following:

Conjecture 1. END OF THE LINE *requires* $2^{\tilde{\Omega}(n)}$ *time*.

It was recently shown that this conjecture implies an identical lower bound for the GCIRCUIT problem [12, 15]: Recall that a standard arithmetic circuit has inputs, outputs, and gates that define arithmetic operations on its lines. In a generalized circuit there are no outputs, and the dependencies between values become cyclical, thereby inducing a constraint satisfaction problem: the variables correspond to the values on the lines, and each gate induces a constraint on the incoming and outgoing lines. The ε -GCIRCUIT is defined to be the problem of finding values on the lines that ε -approximately satisfy all the constraints induced by the arithmetic gates. It was recently shown [26] that ε -GCIRCUIT is PPAD-complete for some small constant ε , via $\tilde{O}(n)$ reductions. Therefore, Conjecture 1 implies that, for some small $\varepsilon > 0$, roughly exponential time is required to find a solution to this problem. There are several known ways to further reduce the ε -GCIRCUIT problem to the two-player Nash equilibrium problem; however, all of these reductions fail to preserve the quality of approximation, and, as we pointed out above, it is likely that no approximation-preserving reductions are possible, as this would imply a quasipolynomial-time algorithm for END OF THE LINE through the algorithm in [23], contradicting Conjecture 1.

In this paper, we identify a plausible new conjecture, a strengthening of Conjecture 1 in a natural and novel direction, which implies that the quasipolynomial approximation algorithm of [23] is optimum. In particular, we define the (ε, δ) -GCIRCUIT to be the problem of finding values for the variables (lines) that ε -approximately satisfy a fraction of at least $(1 - \delta)$ of the constraints (gates).

Conjecture 2. *There exist constants $\varepsilon, \delta > 0$ such that there is a quasilinear reduction from END OF THE LINE to the problem (ε, δ) -GCIRCUIT. Therefore, also assuming Conjecture 1, (ε, δ) -GCIRCUIT requires $2^{\tilde{\Omega}(n)}$ time.*

As we have mentioned, this latter statement, albeit with $\delta = 0$, is known to follow from only Conjecture 1, for some $\varepsilon > 0$ [26].

The Connection to PCP

The (ε, δ) -GCIRCUIT problem is a constraint satisfaction problem: each line/mixed strategy is a real variable, and each gate/player defines a constraint. Since we are considering ε -approximate satisfaction of the constraints, each variable need be represented using only a constant (depending on ε) number of bits. Therefore, a satisfying assignment (truncations of a Nash equilibrium) can be distinguished from an unsatisfying one (one violating the $(\varepsilon - \delta)$ -relaxation) by querying a constant (proportional to $1/\delta$) number of bits of the input, determined at random (but of course not uniformly so). This suggests that Conjecture 2 can be interpreted as a probabilistically checkable proof formulation of GCIRCUIT: informally, it states that “PPAD has a PCP” (that is, a complete problem whose witnesses can be verified by examining, at random, a finite number of bits).

The main result of this paper, explained next, reveals another similarity between the PCP formulation of NP and the (ε, δ) -GCIRCUIT problem: The intractability of the latter problem implies, among several other inapproximability theorems, the strongest possible inapproximability result for the 2-player Nash equilibrium problem, arguably the central open question in the area.

The Main Result

We denote by ε -2NASH the problem of finding an ε -Nash equilibrium in a bimatrix game. Our main result is the following:

Theorem 1. *There is an $\varepsilon > 0$ such that, assuming Conjecture 2, solving ε -2NASH for two-player games with n strategies requires $n^{\tilde{\Omega}(\log n)}$ time.*

Our proof, given in Section 3, employs the technique of “birthday repetition,” pioneered by [1] and used in different contexts by [8], and in particular by [9] to show intractability of a Nash equilibrium-related problem. Starting from a polymatrix game with two strategies per player and in which the player interactions form a cubic bipartite graph with n nodes on each side, the players on both sides are broken into blocks of size \sqrt{n} . The game is simulated by a two-player game, in which each player simulates the nodes in one of the sides of the bipartite graph by choosing a block and a strategy for each node in it; that is, the total number of actions of each of the two players is about $2\sqrt{n}$, and such is the complexity of the reduction (this is necessary if one wishes to rule out better than quasipolynomial algorithms). The interactions between blocks, plus certain particular side games played by the two players, ensure a faithful enough simulation of the original multimatrix game.

Remark: It is not hard to see that a quasipolynomial lower bound follows from a weaker assumption than Conjecture 2, namely that the (ε, δ) -GCIRCUIT problem requires $2^{\tilde{\Omega}(n^\alpha)}$ time, for some $\alpha > \frac{1}{2}$.

In addition to Theorem 1, we prove two other complexity consequences of Conjecture 2, solving certain open problems in the area: An improved inapproximability result for the problem of relative (multiplicative) approximation of the Nash equilibrium first established by Daskalakis [14], and an inapproximability result for the problem of finding a competitive equilibrium with equal incomes and indivisible goods [10, 25] when one seeks to minimize the Gini index of the income distribution.

2 The (ε, δ) -GCIRCUIT and Weak Approximate Nash

Generalized circuits are similar to the standard algebraic circuits, the main difference being that generalized circuits contain cycles, which allow them to verify fixed points of continuous functions.

We restrict the class of generalized circuits to include only a particular list of gates described below. Formally,

Definition 2. [Generalized circuits, [12]] A *generalized circuit* \mathcal{S} is a pair (V, \mathcal{T}) , where V is a set of nodes and \mathcal{T} is a collection of gates. Every gate $T \in \mathcal{T}$ is a 5-tuple $T = G(\zeta \mid v_1, v_2 \mid v)$, in which $G \in \{G_\zeta, G_{\times\zeta}, G_=, G_+, G_-, G_<, G_\vee, G_\wedge, G_\neg\}$ is the type of the gate; $\zeta \in \mathbb{R} \cup \{\text{nil}\}$ is a (optional) real parameter; $v_1, v_2 \in V \cup \{\text{nil}\}$ are the first and second input nodes of the gate (one or both of them may be missing); and $v \in V$ is the output node; no two distinct gates have the same output.

Alternatively, we can think of each gate as a constraint on the values on the incoming and outgoing wires. We are interested in the following constraint satisfaction problem: given a generalized circuit, find an assignment to all the wires that simultaneously satisfies all the gates. When every gate computes a continuous function of the incoming wires, a solution must exist by Brouwer's fixed point theorem.

In particular, we are interested in the approximate version of this CSP, where we must approximately satisfy most of the constraints.

Definition 3. Given a generalized circuit $\mathcal{S} = (V, \mathcal{T})$, we say that an assignment $\mathbf{x}: V \rightarrow [0, 1]$ (ε, δ) -*approximately satisfies* \mathcal{S} if for all but a δ -fraction of the gates, \mathbf{x} satisfies the corresponding constraints:

Gate	Constraint
$G_\zeta(\alpha \mid a)$	$\mathbf{x}[a] = \alpha \pm \varepsilon$
$G_{\times\zeta}(\alpha \mid a \mid b)$	$\mathbf{x}[b] = \alpha \cdot \mathbf{x}[a] \pm \varepsilon$
$G_=(\mid a \mid b)$	$\mathbf{x}[b] = \mathbf{x}[a] \pm \varepsilon$
$G_+(\mid a, b \mid c)$	$\mathbf{x}[c] = \min(\mathbf{x}[a] + \mathbf{x}[b], 1) \pm \varepsilon$
$G_-(\mid a, b \mid c)$	$\mathbf{x}[c] = \max(\mathbf{x}[a] - \mathbf{x}[b], 0) \pm \varepsilon$
$G_<(\mid a, b \mid c)$	$\mathbf{x}[c] = \begin{cases} 1 \pm \varepsilon & \mathbf{x}[a] < \mathbf{x}[b] - \varepsilon \\ 0 \pm \varepsilon & \mathbf{x}[a] > \mathbf{x}[b] + \varepsilon \end{cases}$
$G_\vee(\mid a, b \mid c)$	$\mathbf{x}[c] = \begin{cases} 1 \pm \varepsilon & \mathbf{x}[a] = 1 \pm \varepsilon \text{ or } \mathbf{x}[b] = 1 \pm \varepsilon \\ 0 \pm \varepsilon & \mathbf{x}[a] = 0 \pm \varepsilon \text{ and } \mathbf{x}[b] = 0 \pm \varepsilon \end{cases}$
$G_\wedge(\mid a, b \mid c)$	$\mathbf{x}[c] = \begin{cases} 1 \pm \varepsilon & \mathbf{x}[a] = 1 \pm \varepsilon \text{ and } \mathbf{x}[b] = 1 \pm \varepsilon \\ 0 \pm \varepsilon & \mathbf{x}[a] = 0 \pm \varepsilon \text{ or } \mathbf{x}[b] = 0 \pm \varepsilon \end{cases}$
$G_\neg(\mid a \mid b)$	$\mathbf{x}[b] = \begin{cases} 1 \pm \varepsilon & \mathbf{x}[a] = 0 \pm \varepsilon \\ 0 \pm \varepsilon & \mathbf{x}[a] = 1 \pm \varepsilon \end{cases}$

Given a generalized circuit $\mathcal{S} = (V, \mathcal{T})$, (ε, δ) -GCIRCUIT is the problem of finding an assignment that (ε, δ) -approximately satisfies it.

Definition 4. In an (ε, δ) -weak approximate Nash equilibrium at most a δ -fraction of the players can gain more than ε by deviating.

Theorem 2. Conjecture 2 is equivalent to the following statement: There exist constants $\varepsilon', \delta' > 0$, such that finding an (ε', δ') -weak approximate Nash equilibrium in a polymatrix graphical game with degree 3 and 2 actions per player requires $2^{\tilde{\Omega}(n)}$ time.

Proof. The reduction from (ε, δ) -GCIRCUIT to (ε', δ') -weak approximate Nash equilibrium follows analogously to the reduction in [26] from ε -GCIRCUIT to ε' -approximate Nash equilibrium.

The first step is to reduce (ε, δ) -GCIRCUIT to $(\Theta(\varepsilon^2), \Theta(\varepsilon \cdot \delta))$ -GCIRCUIT with fan-out 2. The naive way to do this is to replace larger fan-outs with a binary tree of $G_=_$ gates. Daskalakis et al. [15] successfully use this method for exponentially small ε , but for constant ε the iterative application of $G_=_$ gates accumulates too much noise ($\Theta(\varepsilon \cdot \log n)$). Instead, we follow [26] and focus on logical gates whose noise does not accumulate. Given an output of a logical gate $(G_<, G_>, G_\wedge, G_\neg)$, we can introduce a binary tree (with even depth) of G_\neg gates that copy the original output instead of having a single gate with a large fan-out. When we begin with an output of an arithmetic gate $(G_\zeta, G_{\times\zeta}, G_+, G_-)$, we first use $\Theta(1/\varepsilon)$ gates of types $G_=_$ and $G_<$ to parse it into an ε -precision unary representation. Then, we copy each bit in the unary representation using a tree of G_\neg gates. Finally, we (approximately) recover the original value from each copy of the unary representation using $\Theta(1/\varepsilon)$ gates of types $G_{\times\zeta}$ and G_+ . Notice that the number of new gates introduced for each gate at the leaves of the binary tree is $\Theta(1/\varepsilon)$. Therefore, if the solution to the new (fan-out 2) instance violates an $O(\delta \cdot \varepsilon)$ -fraction of the gate-constraints, then it induces a solution to the original (arbitrary fan-out) instance that violates at most a δ -fraction of the constraints. (See the full version of [26] for more details.)

We now reduce (ε, δ) -GCIRCUIT with fan-out 2 to (ε', δ') -weak approximate Nash equilibrium. Daskalakis et al. construct, for each gate in the generalized circuit, a game gadget with a few players whose mixed strategies at approximate equilibrium simulate the computation carried by the gate. In other words, every approximate equilibrium of the gate gadget induces an approximately satisfying assignment to the corresponding input and output lines in the generalized circuit. See Lemma 1 for an example of such a gadget. The gadgets are concatenated by identifying the “output player” of one gadget with the “input player(s)” of the next gadget(s) in the generalized circuit. Since each gadget is composed of only a constant number of players, if all but a δ' -fraction of the players play approximately at equilibrium, all but a $\delta = \Theta(\delta')$ -fraction of the gates are violated by the induced assignment to the generalized circuit.

The above construction suffices to show hardness for an $(\hat{\varepsilon}, \delta')$ -weak well-supported Nash equilibrium, i.e. one for which all but a δ' -fraction of the players, every action in the support is $\hat{\varepsilon}$ -approximately best response. Finally, for constant degree graphical games, it is not hard to derive an $(\hat{\varepsilon}, \delta')$ -weak well-supported Nash equilibrium from any (ε', δ') -weak approximate Nash equilibrium for $\varepsilon' = \Theta(\hat{\varepsilon}^2)$ [26, Lemma 4].

In the other direction, from (ε', δ') -weak approximate Nash equilibrium to (ε, δ) -GCIRCUIT, it suffices to construct a generalized circuit that computes an ε' -approximate best response of each player (except a δ' -fraction). Notice that in a graphical game with a constant number of actions per player, given a profile of mixed strategies, one only needs a constant number of arithmetic gates to compute a best response of any player. In particular, to compute an ε -approximate best response it suffices to use gates of finite precision $\varepsilon' = \Theta(\varepsilon)$. Finally, since the number of gates per player is constant, if at most a δ -fraction of the gates are ε -unsatisfied, they can appear in the best-response computation for at most a $\delta' = \Theta(\delta)$ -fraction of the players. \square

Lemma 1. ($G_{\times\zeta}$ gadget, [15])

Let v_1, v_2 , and w be players in a graphical game, and suppose that the payoffs of v_2 and w are as follows.

		w plays 0	w plays 1
Payoff for v_2 :	v_2 plays 0	0	1
	v_2 plays 1	1	0

Payoffs for w :

game with v_1 :		$v_1 \text{ plays } 0$	$v_1 \text{ plays } 1$
	$w \text{ plays } 0$	0	ζ
	$w \text{ plays } 1$	0	0
game with v_2 :		$v_2 \text{ plays } 0$	$v_2 \text{ plays } 1$
	$w \text{ plays } 0$	0	0
	$w \text{ plays } 1$	0	1

Then, in every ε -NE $\mathbf{p}[v_2] = \min(\zeta \mathbf{p}[v_1], 1) \pm \varepsilon$, where $\mathbf{p}[u]$ denotes the probability that u assigns to strategy 1.

Proof. (Sketch) If $\mathbf{p}[v_2] > \zeta \mathbf{p}[v_1] + \varepsilon$, then in every ε -NE $\mathbf{p}[w] = 1$, which contradicts $\mathbf{p}[v_2] > \varepsilon$. Similarly, if $\mathbf{p}[v_2] < \min(\zeta \mathbf{p}[v_1], 1) - \varepsilon$, then $\mathbf{p}[w] = 0$, which yields a contradiction when $\mathbf{p}[v_2] < 1 - \varepsilon$. \square

3 Proof of Theorem 1

3.1 Overview

The proof is a reduction from weak approximation of Nash equilibrium in a polymatrix game (recall Theorem 2) to the two-player problem. The two players simultaneously play three games: the main game, which is the heart of the reduction; and two games based on a construction due to Althofer [2], which impose structural properties of any approximate Nash equilibrium (interestingly, the oblivious lower bound of [16] uses the same game). The final payoff of a player is the sum of payoffs in all three games. For convenience of notation, the payoffs in each game will be in $[0, 1]$; in order to normalize the payoffs in the final game to $[0, 1]$ one should multiply by $1/3$ the payoffs in all three games.

Main game

We let each of the two players “control” the vertices on one side of the bipartite graphical game (we henceforth use *players* only for the players in the bimatrix game, and refer to the players in the graphical polymatrix game as *vertices*). For ease of notation, we assume wlog that each player controls an equal number of vertices, n .

We partition the vertices of each player into n/k disjoint subsets of size at most $2k = 2\sqrt{n}$, such that every two subsets share at most 18 edges. By Lemma 6, we can efficiently find such a partition. Let $(S_1, \dots, S_{n/k})$ and $(T_1, \dots, T_{n/k})$ denote the partitions of the respective players. Each action of the players corresponds to a choice of a subset (out of n/k subsets in the partition), and a choice of an action for each vertex in the subset (out of at most 2^{2k} vectors of actions). Together, the main game has $(\frac{n}{k} \cdot 2^{2k}) \times (\frac{n}{k} \cdot 2^{2k})$ action profiles. When players choose actions $(S_i, \vec{\alpha}_{S_i})$ and $(T_j, \vec{\beta}_{T_j})$, the payoff of the row player is the sum of the payoffs over all shared edges of S_i and T_j , when they play the respective strategies from $\vec{\alpha}_{S_i}$ and $\vec{\beta}_{T_j}$ (here we use the polymatrix structure of the payoffs in the graphical game; the payoffs are defined over edges and therefore payoffs of a certain vertex can be defined even though not necessarily all its neighbours are playing). Finally, we normalize by $\lambda/18$, where $\lambda = \Theta(\delta^2)$ is a small constant that satisfies $\lambda > \varepsilon'$. Similarly the payoff of the second player is derived from the payoffs of the vertices in T_j .

Althofer's games

Althofer's game [2] is an asymmetric hide-and-seek win-lose game over l locations. The hider chooses a location $i \in [l]$, the seeker chooses a subset $B \subset [l]$ of size $|B| = l/2$. The hider wins if $i \notin B$, the seeker wins if $i \in B$. Namely, the payoff function is given by

$$u_1(i, B) = 1 - u_2(i, B) = \begin{cases} 1 & \text{if } i \notin B \\ 0 & \text{if } i \in B. \end{cases}$$

In Althofer's game each player can guarantee $\frac{1}{2}$ by uniform play, therefore the value of the game is $\frac{1}{2}$. Althofer's game enjoys the following strong property: in every ε -approximate Nash equilibrium, the hider must play a mixed strategy that is $O(\varepsilon)$ -close to the uniform distribution in total variation distance; see [16].

In our game each player plays two (unrelated) Althofer's games, one game as a hider, and one as a seeker. When player 1 is a hider, the locations are the sets $S_1, \dots, S_{n/k}$ (i.e., $l = n/k$). The action from the main game $(S_i, \vec{\alpha}_{S_i})$ determines the location, where for the purposes of Althofer's game we ignore $\vec{\alpha}_{S_i}$. When player 1 is a seeker, the locations are the sets $T_1, \dots, T_{n/k}$, and player 1 chooses an action to play in this game, independent of $(S_i, \vec{\alpha}_{S_i})$. The total number of actions for player 1 is $\left(\frac{n}{k} \cdot 2^{2k}\right) \cdot \binom{n/k}{n/(2k)} = 2^{O(\sqrt{n})}$. Similarly, for player 2, who is playing as a hider over $T_1, \dots, T_{n/k}$, and as a seeker over $S_1, \dots, S_{n/k}$.

3.2 Structure of an equilibrium

For a mixed action x of player 1 we denote by $x(S_i)$ the total probability that the player chooses to control the vertices S_i ; i.e.,

$$x(S_i) = \sum_{\vec{\alpha} \in 2^{[k]}, B \subset \{T_1, \dots, T_{n/k}\}, |B| = \frac{n}{2k}} x(((S_i, \vec{\alpha}), B)).$$

Lemma 2. *If (x, y) is an ε' -Nash equilibrium, then $\sum_{i \in [n/k]} |x(S_i) - \frac{k}{n}| = O(\lambda)$.*

In the proof we use the following Lemma from [16]:

Lemma 3 ([16]). *Let $\{a_i\}_{i \in [l]}$ be real numbers satisfying the following properties for some $\theta > 0$: (1) $a_1 \geq a_2 \geq \dots \geq a_l$; (2) $\sum_{i \in [l]} a_i = 0$; and (3) $\sum_{i \in [l/2]} a_i \leq \theta$. Then $\sum_{i \in [l]} |a_i| \leq 4\theta$.*

Proof of Lemma 2. In order apply Lemma 3, we denote $a_i = x(S_i) - \frac{k}{n}$ and we assume wlog that $x(S_1) \geq x(S_2) \geq \dots \geq x(S_{n/k})$. Then the first two conditions hold. Regarding the third condition, we argue that $\sum_{i \in [l/2]} x(S_i) \leq 3\lambda$ this will complete the proof.

Player 1 can guarantee a payoff of $1/2$ in the sum of Althofer's games as a hider and the main game, by playing uniformly and choosing an arbitrary actions for the controlled vertices (for instance, $\vec{\alpha}_{S_i} = \vec{0}$). Assume by contradiction that $\sum_{i \in [l/2]} x(S_i) > 3\lambda$. Since Player 2 is ε' -best replying, his payoff in the Althofer game (as a seeker) is at least $\frac{1}{2} + 3\lambda - \varepsilon'$ (because he can get $1/2 + 3\lambda$ by choosing the set $[l/2]$). Therefore, Player 1's payoff in the Althofer game (as a hider) is at most $\frac{1}{2} - 3\lambda + \varepsilon'$. If we add to it his payoffs in the main game, then his payoff is at most $\frac{1}{2} - 2\lambda + \varepsilon'$. Therefore, player 1 can increase his payoff by $2\lambda - \varepsilon' > \varepsilon'$ by deviating to the uniform distribution over the locations (as a hider), and maintaining his mixed action as a seeker. \square

3.3 Completing the proof

Any mixed strategy x of player 1 in the bimatrix game induces a mixed strategy of all the vertices in $\cup_i S_i$ in the obvious way. Vertex $s \in S_i$ plays the action 1 with probability p where p is the conditional probability

$$Pr_x(\vec{\alpha}_{S_i}(s) = 1 | \text{Player 1 controls } S_i).$$

If the event “Player 1 controls S_i ” occurs with probability 0, then we define (arbitrarily) that $p = 1$. Similarly, any mixed strategy y of player 2 induces a mixed strategy of all the vertices in $\cup_i T_i$.

We claim that if (x, y) is an ε' -approximate Nash equilibrium of the bimatrix game, then the induced mixed-strategies profile is an (ε, δ) -approximate Nash equilibrium of the original graphical game.

By Lemma 2 and Markov’s inequality, for a $(1 - O(\sqrt{\lambda}))$ -fraction of the subsets, the row player distributes within a $(1 \pm O(\sqrt{\lambda}))$ -factor of the correct weight (k/n) ; i.e. $x(S_i) \in \frac{k}{n} \cdot [1 - O(\sqrt{\lambda}), 1 + O(\sqrt{\lambda})]$. Let us restrict our attention only to those subsets, and call them *good*. We say that a vertex is *good* if it and all of its neighbors belong to good subsets. Since the game graph is of bounded degree, we again have that a $(1 - O(\sqrt{\lambda}))$ -fraction of the vertices is good.

Consider any good vertex whose induced mixed strategy is not ε -optimal in the original game given that the rest of the vertices also play according to their induced strategies. Then, changing its strategy in the bimatrix game (while leaving all other marginals the same), would increase the payoff of its player by at least

$$(1 - O(\sqrt{\lambda}))^2 \cdot \left(\frac{k}{n}\right)^2 \cdot (1 - O(\sqrt{\lambda})) \cdot \varepsilon \cdot \frac{\lambda}{18} = \Omega(\varepsilon \lambda / n),$$

where the $(1 - O(\sqrt{\lambda}))^2 \cdot (\frac{k}{n})^2$ term corresponds to the probabilities that the subsets corresponding to both the vertex and any of its neighbors are played; the $(1 - O(\sqrt{\lambda})) \cdot \varepsilon$ term corresponds to the improvement of the vertex in the bimatrix game, where instead of summing the payoff in all three edges (as in the graphical game) these three edges have weights $\{1 + \gamma_j\}_{j=1,2,3}$ for $|\gamma_j| = O(\sqrt{\lambda})$; finally $\lambda/18$ is the normalization. The right side follows by plugging in $k = \sqrt{n}$.

If (by a contradiction) a δ -fraction of the vertices have ε -improvement strategies; then one of the players has $(\delta/2 - O(\sqrt{\lambda})) \cdot n$ good vertices with an ε -improvement¹. This player can benefit $\Omega(\varepsilon \cdot \lambda / n)$ from a deviation of each vertex. So his total improvement from all deviations simultaneously is $\Omega(\varepsilon \cdot \delta \cdot \lambda)$ - which is impossible when (x, y) is an ε' -approximate Nash equilibrium for $\varepsilon' = \Theta(\varepsilon \cdot \delta \cdot \lambda)$.

4 Implications for Relative ε -approximate Nash

Daskalakis [14] defines a notion of relative (sometimes also called *multiplicative* [20], as opposed to the more standard additive) ε -Nash equilibrium, and proves that in two-player games with payoffs in $[-1, 1]$, finding a relative ε -well supported Nash equilibrium is PPAD-complete. In particular, he

¹Here we use our choice of $\lambda = \Theta(\delta^2)$, which guarantees that $\delta/2 - O(\sqrt{\lambda}) = \Omega(\delta)$

concludes that the quasi-polynomial algorithm of Lipton et al. [23] cannot achieve this notion of approximate equilibrium.

This result has two caveats: (1) Through the use of positive and negative payoffs, the gain from deviation is large compared to the expected payoff, only because the latter is small due to cancellation of positive and negative payoffs. Namely, the gain from deviation may be very small compared to the expected magnitude of the payoff. (2) It only applies to the more restrictive notion (thus rendering the hardness result weaker) of *well-supported* approximate equilibrium; i.e. an equilibrium where every action in the support has to be approximately optimal. Showing PPAD-hardness for both positive payoffs and general (non well-supported) approximate equilibrium were left as open questions in [14]. Recently, the first question was settled in [26] where it was shown that finding a relative ε -well-supported Nash equilibrium with positive payoffs is indeed PPAD-complete.

Here, assuming Conjecture 2, we settle the second question: we show that finding any relative ε -approximate Nash equilibrium is PPAD-complete. Furthermore, in our hard instance the row player has only positive payoffs and the column player has only negative payoffs, and so there is no cancellation of payoffs as in the construction of [14].

Theorem 3. *Assuming Conjecture 2 there exists a constant $\varepsilon' = \Theta(\varepsilon \cdot \delta^3) > 0$ such that finding a relative ε' -approximate Nash equilibrium (ANE) in a bimatrix game where the row player's payoffs are non-negative and the column player's payoffs are non-positive is PPAD-complete.*

Proving a similar theorem when both players have positive payoffs remains an interesting open question. In fact, we do not know of any instances with positive payoffs where all ε -approximate Nash equilibria must have large (e.g., linear, or even super-logarithmic) support.

Proof. We reduce from the problem of finding an (ε, δ) -Nash in a bipartite, degree 3, polymatrix game with two actions per player. We construct a main game where the bimatrix game players (henceforth just *players*) control the nodes of the polymatrix game, and two side games that guarantee that each player randomizes (approximately) uniformly over all her nodes.

Main game We let the row player “control” the nodes on one side of the bipartite game graph, and let the column player control the nodes on the other side. Namely, let n be the number of nodes on each side of the graph; each player has $2n$ actions, each corresponding to a choice of node and an action for that node. When the players play strategies that correspond to adjacent nodes in the graphs, they receive the payoffs of the corresponding nodes, scaled (by a small positive constant $\eta = O(\delta^2)$) and shifted to fit in the intervals: $[1, 1 + \eta]$ for the row player, and $[-(1 + \eta), -1]$ for the column player. If the nodes played do not share an edge in the bipartite game graph, the utility for both players is zero. Notice that if either player chooses her node uniformly at random (and chooses the action for that node arbitrarily), then the expected payoff is in $\frac{3}{n} \cdot [1, 1 + \eta]$ for the row player, and in $\frac{3}{n} \cdot [-(1 + \eta), -1]$ for the column player.

Side games We also let the players play two hide-and-seek win-lose zero-sum games over a space of n actions. In both games, the row player is chasing the column player. In each game, if they pick the same strategy, the row player receives payoff 1 (and the column player receives -1); otherwise the payoffs are 0. Finally, in the first side game we identify the row player's strategies with her choice of nodes in the main game. Namely, if she plays node i in the main game and the column player chose i in the first side game, then her payoff from this side game is 1. Similarly, we identify the column player's strategies in the second game with her choice of nodes in the main game.

We proceed by showing that in every relative ε' -ANE, the row player's utility is approximately $5/n$, and the column player's utility is approximately $-5/n$ (Lemma 4); then we show that in every relative ε' -ANE, both players randomize approximately uniformly over their nodes (Lemma 5); finally we use these two observations to complete the proof (Subsection 4.2). \square

4.1 Value and structure of a relative ε' -ANE

Given mixed strategies (x, y) , we let $x(i)$ denote the total probability that the row player assigns to node i , and analogously for $y(j)$. We also let $x^*(y)$ and $y^*(x)$ denote the corresponding best responses of each player. Finally, let $U_R(x; y)$ and $U_C(y; x)$ denote the expected payoffs for the row and column players, respectively.

Lemma 4. *If (x, y) is a relative ε' -ANE, then*

$$(1 - \varepsilon') \cdot \frac{5}{n} \leq U_R(x; y) \leq (1 + \eta)(1 + \varepsilon') \cdot \frac{5 + 3\eta}{n} \quad (1)$$

and

$$-(1 - \varepsilon')(1 - \eta) \cdot \frac{5}{n} \geq U_C(y; x) \geq -(1 + \varepsilon') \cdot \frac{5 + 3\eta}{n} \quad (2)$$

Proof. Observe that the main game is relative η -approximately zero-sum; i.e. for any pure strategy profile (x', y') the expected utilities $U_R^{\text{main}}(x'; y')$, $U_C^{\text{main}}(y'; x')$ of the row and column player, respectively, satisfy:

$$\left| U_R^{\text{main}}(x'; y') + U_C^{\text{main}}(y'; x') \right| \leq \eta \cdot \min \left\{ U_R^{\text{main}}(x'; y'), -U_C^{\text{main}}(y'; x') \right\} \quad (3)$$

By linearity of expectation and convexity of the absolute value function, this continues to hold when (x', y') are mixed strategies. Furthermore, since the side games are exactly zero-sum and have the same signs as the main game, the same follows for the payoffs from the entire game:

$$\left| U_R(x'; y') + U_C(y'; x') \right| \leq \eta \cdot \min \left\{ U_R(x'; y'), -U_C(y'; x') \right\} \quad (4)$$

In particular, the above inequality holds for (x, y) . Thus, the upper bounds in (1) and (2) follow from the lower bounds in (2) and (1), respectively.

Finally, $(1 - \varepsilon') \cdot \frac{5}{n} \leq U_R(x; y)$ in every relative ε' -ANE, because the row player can guarantee an expected payoff of at least $\frac{5}{n}$ by randomizing uniformly over all her strategies. Similarly, $U_C(y; x) \geq -(1 + \varepsilon') \left(\frac{3(1+\eta)+2}{n} \right)$ because the column player can guarantee a payoff of at least $-\frac{3(1+\eta)+2}{n}$ by randomizing uniformly over all her strategies. \square

Lemma 5. *There exists a constant $\lambda = \Theta(\delta^2)$ such that for every (x, y) relative ε' -ANE, $\sum |x(i) - \frac{1}{n}| \leq \lambda$ and $\sum |y(j) - \frac{1}{n}| \leq \lambda$.*

Proof. Assume by contradiction that $\sum |x(i) - \frac{1}{n}| > \lambda$, then in particular there must exist an $i \in [n]$ such that $x(i) < (1 - \lambda/2)/n$. When the column player chooses her strategy uniformly at random in the main game and in the second side game and plays strategy i in the first side game, her expected payoff is at least

$$U_C(y^*(x); x) \geq -\frac{3}{n}(1 + \eta) - \frac{1}{n} - x(i).$$

Therefore, in any relative ε' -ANE her expected utility is at least

$$U_C(y; x) \geq (1 + \varepsilon') \cdot \left(\frac{-5 - 3\eta + \lambda/2}{n} \right),$$

Which contradicts the upper bound in (2) when we take λ sufficiently large, e.g. $\lambda = 10\eta$.

Similarly, if $\sum |y(j) - \frac{1}{n}| > \lambda$, there must exist an $j \in [n]$ such that $y(j) > (1 + \lambda/2)/n$. Therefore, the row player can guarantee a payoff of at least

$$U_R(x^*(y); y) \geq \frac{4}{n} + y(j).$$

Thus by relative ε' -ANE,

$$U_R(x; y) \geq (1 - \varepsilon') \cdot \left(\frac{5 + \lambda/2}{n} \right),$$

which contradicts the upper bound in (1). \square

4.2 Completing the proof of Theorem 3

Now, given a relative ε' -ANE (x, y) , we can take, for each node i , the mixed strategy induced by the probabilities $x(i : 1)/x(i)$ and $x(i : 2)/x(i)$ that the row player assigns to each action (respectively, $y(j : 1)/y(j)$ and $y(j : 2)/y(j)$ assigned by the column player). We claim that this strategy profile is an (ε, δ) -approximate equilibrium for the polymatrix game. Assume by contradiction that this is not the case.

By Lemma 5 and Markov's inequality, a $(1 - O(\sqrt{\lambda}))$ -fraction of the nodes are played within $\frac{\sqrt{\lambda}}{n}$ of the correct probability $1/n$. For a node i controlled by the row player, we say that it is *good* if $|x(i) - 1/n| \leq \frac{\sqrt{\lambda}}{n}$ and $|y(j) - 1/n| \leq \frac{\sqrt{\lambda}}{n} \forall j \in \mathcal{N}(i)$, and analogously for column player's nodes. Since the graph has bounded degree, a $(1 - O(\sqrt{\lambda}))$ -fraction of the nodes are good.

Let i be any good node who has an ε -improving deviation from her induced strategy in the polymatrix game. If the player who controls i makes the corresponding deviation in the two-player game, she increases her expected payoff by at least $\varepsilon \cdot \eta \cdot \frac{1}{n^2} \cdot (1 - O(\sqrt{\lambda}))$. (We multiply by η to account for the scaling; by $1/n^2$ for the probability that this player plays node i and the other player plays a neighbor of i ; and by $(1 - O(\sqrt{\lambda}))$ to correct for the deviation from $1/n$ in the latter probabilities.) By our assumption that the induced strategy profile is not an (ε, δ) -approximate equilibrium for the polymatrix game, at least one of the players has at least $(\delta - O(\sqrt{\lambda}))n$ good nodes with ε -improving deviations. Summing her gains from those deviations, we get that this player has can improve her expected payoff by at least $\left[\varepsilon \cdot \eta \cdot (1 - O(\sqrt{\lambda})) \right] (\delta - O(\sqrt{\lambda})) / n$.

However, this is a contradiction since it follows from Lemma 4, that (x, y) is also an (additive) $\varepsilon' \cdot (\frac{\delta}{n})$ -ANE. \square

5 Implications for Fairness Mechanisms

Competitive equilibrium from equal incomes (CEEI) is a well-known fair allocation mechanism [18, 29, 27]: We give all agents a unit of money, and price the goods in such a way that the market clears. It is also well-known that when goods are indivisible or utilities are non-linear, an equilibrium may not exist. However, Budish [10] proves that an approximate CEEI still exists.

This concept of approximate equilibrium is used in practical system for allocating seats in courses to business school students [11].

Budish [10] measures the proximity to a perfect CEEI via two parameters: a solution is an (α, β) -CEEI if the clearing error of the competitive equilibrium is less than α , and all the incomes are between 1 and $1 + \beta$. Budish shows that an (α, β) always exists for some favorable $\alpha = \alpha^*$, and any $\beta > 0$. Recently, [25] showed a reduction from ε -GCIRCUIT with fan-out 2, to the problem of finding an $(\alpha^*, \Theta(\varepsilon/\log(1/\varepsilon)))$ -CEEI. In particular, when combined with the results of [26], this implies that it is PPAD-complete to find an (α^*, β) -CEEI for some constant $\beta > 0$.

The β parameter used in Budish’s formulation is an imperfect way of measuring income inequalities, as it may be set by a single outlier. Perhaps the best known, and most widely used, measure of income inequality is the *Gini index* (e.g. [19, 13]) (see the Appendix for the definition). In fact, the Gini index is used to assess the performance of the class seat assignment system used in practice [11].

Theorem 4 (Informal). *Assuming Conjecture 2, finding an income assignment and prices with low market clearing error and near-optimal Gini index is PPAD-complete.*

See the appendix for the precise definitions and statement of the result.

6 Discussion

The purpose of this paper is to showcase an important open problem:

reduce END OF THE LINE **to** (ε, δ) -GCIRCUIT

— i.e., show that Conjecture 1 implies Conjecture 2. As we mentioned, such a reduction would imply a “PCP for PPAD”; i.e. it would imply that there exists a probabilistically checkable proof for a PPAD-complete problem.

The equivalent statement for the class of NP-complete problem, is the celebrated PCP Theorem [4, 3]. What can we learn from our experience with constructing PCP’s for NP? First of all, we have many different constructions of PCPs for NP, and this may be seen as circumstantial evidence that constructing a PCP for PPAD is possible. Equally important, there are several constructions of PCPs for NP that are near-linear (e.g. [6, 7, 17, 24]); notice that for the application to ε -Nash in bimatrix games we need the proof length to be sub-quadratic.

The next question one should ask is, whether we can adapt the techniques used in the proofs of the PCP Theorem to the class of PPAD-complete problems. We briefly and informally sketch some of our thoughts on the matter. All the proofs of the PCP Theorems that we are aware of (including [4, 3, 17]) compose an *inner verifier* with an *outer verifier*. The outer verifier in Dinur’s proof [17] is combinatorial in nature, and it seems plausible that the same or similar techniques could be modified to fit the generalized circuit graph. Unfortunately, all inner verifiers that we know are discrete in nature, whereas our characterization of PPAD with GCIRCUIT (or Nash) is inherently based on continuous constraints. Thus the following interesting (and somewhat open-ended) questions arise: Can we find a characterization of PPAD via a constraint satisfaction problem whose constraints are discrete in nature? Alternatively, can one construct an inner (non-efficient) verifier using the constraints specified in the definition of the GCIRCUIT problem?

Acknowledgments: Many thanks to Boaz Barak, Paul Cristiano, Muli Safra, and Madhu Sudan for inspiring discussions about probabilistic checkable proofs.

References

- [1] Scott Aaronson, Russell Impagliazzo, and Dana Moshkovitz. Am with multiple merlins. In *Computational Complexity (CCC), 2014 IEEE 29th Conference on*, pages 44–55. IEEE, 2014.
- [2] Ingo Althofer. On sparse approximations to randomized strategies and convex combinations. 1993.
- [3] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [4] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.
- [5] Paul Beame, Stephen A. Cook, Jeff Edmonds, Russell Impagliazzo, and Toniann Pitassi. The relative complexity of NP search problems. *J. Comput. Syst. Sci.*, 57(1):3–19, 1998.
- [6] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil P. Vadhan. Robust pcps of proximity, shorter pcps, and applications to coding. *SIAM J. Comput.*, 36(4):889–974, 2006.
- [7] Eli Ben-Sasson and Madhu Sudan. Robust locally testable codes and products of codes. *Random Struct. Algorithms*, 28(4):387–402, 2006.
- [8] Mark Braverman, Young Kun-Ko, Aviad Rubinstein, and Omri Weinstein. ETH Hardness for Densest- k -Subgraph with Perfect Completeness. In submission.
- [9] Mark Braverman, Young Kun-Ko, and Omri Weinstein. Approximating the best nash equilibrium in $n^{o(\log n)}$ -time breaks the exponential time hypothesis. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 970–982, 2015.
- [10] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061 – 1103, 2011.
- [11] Eric Budish, Gerard P. Cachon, Judd Kessler, and Abraham Othman. Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. Working paper, 2014.
- [12] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *J. ACM*, 56(3), 2009.
- [13] Frank Cowell. *Measuring inequality*. Oxford University Press, 2011.
- [14] Constantinos Daskalakis. On the complexity of approximating a nash equilibrium. *ACM Transactions on Algorithms*, 9(3):23, 2013.
- [15] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. *Commun. ACM*, 52(2):89–97, 2009.
- [16] Constantinos Daskalakis and Christos H Papadimitriou. On oblivious ptas’s for nash equilibrium. *arXiv preprint arXiv:1102.2280*, 2011.
- [17] Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):12, 2007.

- [18] D.K. Foley. Resource Allocation and the Public Sector. *Yale Economic Essays*, 7(1):45–98, 1967.
- [19] Corrado Gini. Variabilità e mutabilità. *Reprinted in Memorie di metodologica statistica (Ed. Pizetti E, Salvemini, T).* Rome: Libreria Eredi Virgilio Veschi, 1, 1912.
- [20] Sébastien Hémon, Michel de Rougemont, and Miklos Santha. Approximate Nash Equilibria for Multi-player Games. In *SAGT*, pages 267–278, 2008.
- [21] Michael D. Hirsch, Christos H. Papadimitriou, and Stephen A. Vavasis. Exponential lower bounds for finding brouwer fix points. *J. Complexity*, 5(4):379–416, 1989.
- [22] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
- [23] Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *EC*, pages 36–41, 2003.
- [24] Dana Moshkovitz and Ran Raz. Sub-constant error probabilistically checkable proof of almost-linear size. *Computational Complexity*, 19(3):367–422, 2010.
- [25] Abraham Othman, Christos H. Papadimitriou, and Aviad Rubinstein. The complexity of fairness through equilibrium. In *EC*, pages 209–226, 2014.
- [26] Aviad Rubinstein. Inapproximability of nash equilibrium. In *STOC*, 2015. To appear.
- [27] W. Thomson and H.R. Varian. Theories of justice based on symmetry. *Social Goals and Social Organizations: Essays in Memory of Elisha Pazner*, 1985.
- [28] Haralampos Tsaknakis and Paul G. Spirakis. An optimization approach for approximate nash equilibria. *Internet Mathematics*, 5(4):365–382, 2008.
- [29] H. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.

A Finding a good partition

Lemma 6. *Let $G = (U, V, E)$ be a bipartite d -regular graph with $n = |U| = |V|$. Then we can efficiently find partitions $S_1, \dots, S_{n/k}$ and $T_1, \dots, T_{n/k}$ of U and V , respectively, to disjoint subsets, such that each subset has size at most $2k = 2\sqrt{n}$, and:*

$$\forall i, j \in [n/k] \quad |(S_i \times T_j) \cap E| < 2d^2k^2/n.$$

Proof. Let $S_1, \dots, S_{n/k}$ be an arbitrary partition of U into disjoint subsets of size exactly k . We inductively place the vertices of V into subsets $T_1, \dots, T_{n/k}$, while maintaining the desiderata that each subset T_j is of size at most $2k$, and for every i, j the number of edges from S_i to T_j is at most $2d^2k^2/n$.

It is left to show that for any partial partition of V , there is a subset T_j into which we can place the next vertex v . In expectation, every subset T_j has less than k vertices. Therefore by Markov’s inequality, less than half of the subsets have $2k$ vertices or more. v has neighbors in at most d subsets S_i . Recall that the S_i ’s are of size exactly k . Thus each S_i with a neighbor of v has, in expectation over j , less than dk^2/n neighbors in T_j . Using Markov’s inequality again, less than a $1/(2d)$ -fraction of the subsets T_j contain $2d^2k^2/n$ neighbors of S_i . In total, we lose less than

half the subsets for the size desideratum, and less than $1/(2d)$ -fraction of the subsets for each S_i containing a neighbor of v . Therefore there always remains at least one subset T_j to which we can add v . \square

B The course allocation problem

Even though the approximate CEEI and the existence theorem in [10] are applicable to a broad range of allocation problems, we shall describe our results in the language of the course allocation problem.

We are given a set of M courses with integer capacities (the supply) $(q_j)_{j=1}^M$, and a set of N students, where each student i has a set $\Psi_i \subseteq 2^M$ of permissible course bundles, with each bundle containing at most $k \leq M$ courses. The set Ψ_i encodes both scheduling constraints (e.g., courses that meet at the same time) and any constraints specific to student i (e.g. prerequisites).

Each student i has a strict ordering over her permissible schedules, denoted by \preccurlyeq_i . We allow arbitrarily complex preferences — in particular, students may regard courses as substitutes or complements. More formally:

Definition 5. Course Allocation Problem The input to a course allocation problem consists of:

- For each student i a set of course bundles $(\Psi_i)_{i=1}^N$.
- The students' reported preferences, $(\preccurlyeq_i)_{i=1}^N$,
- The course capacities, $(q_j)_{j=1}^M$, and

The output to a course allocation problem consists of:

- Prices for each course $(p_j^*)_{j=1}^M$,
- Allocations for each student $(x_i^*)_{i=1}^N$, and
- Budgets for each student $(b_i^*)_{i=1}^N$.

The quality of an allocation is evaluated based on its proximity to market clearing and income equality. As for the definition of market clearing error, it suffices for our purposes to require that no course is over-subscribed or under-subscribed by more than α^* students (we refer the curious reader to [10] for the precise definition). The Gini coefficient, which measures income inequality, is discussed in the following subsection.

B.1 The Gini coefficient

Definition 6 (Gini index). Given a distribution of incomes D , the *Lorenz curve* plots, for each $x \in [0, 1]$, the cumulative wealth owned by the bottom x -fraction of the population. Let $F_D^{-1}(x) = \sup \{y: \Pr_{z \sim D}[z \leq y] \leq x\}$, and define $L_D(x) = (\int_0^x F_D^{-1}(x)) / (\int_0^1 F_D^{-1}(x))$. Then the Lorenz curve is the graph $(x, L_D(x))$, for $x \in [0, 1]$. Notice that $F_D^{-1}(x)$ is monotonically non-decreasing, so the Lorenz curve is convex.

The *Gini index* is defined as the ratio of the area between the (x, x) line and the Lorenz curve (by the convexity, $x \geq L_D(x) \forall x \in [0, 1]$), divided by the entire area under the (x, x) line (the latter is always $1/2$): $G_D = 2 \int_0^1 (x - L_D(x)) dx = 1 - 2 \int_0^1 L_D(x) dx$

In general, a smaller Gini index corresponds to a more equal distribution of wealth. For example, when all incomes are exactly equal, the Lorenz curve is exactly equal to the (x, x) line, and the Gini index is 0. On the other extreme, when one person has all the wealth, the area under the Lorenz curve goes to 0 as the population size increases, and the Gini index approaches 1.

B.2 Intractability of approximate CEEI with near-optimal Gini index

Theorem (Theorem 4, formal). *Assuming Conjecture 2, there exists some constant $\gamma > 0$ such that finding an allocation with market clearing error α^* and Gini coefficient γ is PPAD-complete.*

The rest of this section is devoted to sketching a proof of Theorem 4. In the next subsection, we briefly outline the reduction of [25] from generalized circuits with fan-out 2 to approximate CEEI. (Recall that in Section 2 we show how to convert any generalized circuit to one with fan-out 2.) Then, we show that after normalizing the median budget to $1 + \varepsilon'/2$, in any allocation which does not correspond to a valid solution to the (ε, δ) -GCIRCUIT instance, a δ' -fraction of the students have budgets either at most 1 or at least $1 + \varepsilon'$, (for some constants $\varepsilon' = \Theta(\varepsilon/\log(1/\varepsilon))$ and $\delta' = \Theta(\delta/\log(1/\varepsilon))$). Then, the proof is complete with the following lemma:

Lemma 7. *Let the median income be $1 + \varepsilon'/2$, and suppose that a δ' -fraction of the population has income at most 1 (resp. at least $1 + \varepsilon'$). Then the Gini index is at least $\gamma = \Theta(\delta' \cdot \varepsilon') = \Theta(\delta \cdot \varepsilon/\log^2(1/\varepsilon))$.*

Proof. The total income of the poorer half of the population is at most $(1/2 - \delta')(1 + \varepsilon'/2) + \delta' = (1/2)(1 + \varepsilon'/2) - \delta' \cdot \varepsilon'/2$, whereas the richer half of the population has a total income of at least $(1/2)(1 + \varepsilon'/2)$. Therefore, the Lorenz curve's value at $1/2$ is bounded by

$$L(1/2) \leq \frac{(1/2)(1 + \varepsilon'/2) - \delta' \cdot \varepsilon'/2}{(1 + \varepsilon'/2) - \delta' \cdot \varepsilon'/2} = 1/2 - \Theta(\delta' \cdot \varepsilon').$$

We can now use elementary geometry to bound the area under the Lorenz curve:

$$\begin{aligned} \int_0^1 L(x)dx &= \int_0^{1/2} L(x)dx + \int_{1/2}^1 L(x)dx \\ &\leq \frac{1}{4} \left[0 + (1/2 - \Theta(\delta' \cdot \varepsilon')) \right] + \frac{1}{4} \left[(1/2 - \Theta(\delta' \cdot \varepsilon')) + 1 \right] \\ &= 1/2 - \Theta(\delta' \cdot \varepsilon'). \end{aligned}$$

Therefore the Gini index is at least $\Theta(\delta' \cdot \varepsilon')$. A similar argument works when a δ' -fraction of the population has income at least $1 + \varepsilon'$. □

B.3 From generalized circuits to Course Allocation

[25] reduce generalized circuits to Course Allocation, by constructing a gadget for each gate of the generalized circuit. (In fact, a few gadgets per gate are required to handle a subtle issue that [25] call "course-size amplification".) We provide the gadget for the NOT gate as an example, and then describe the properties of the [25]'s reduction that we need to complete the proof.

We henceforth normalize the prices and budgets in every assignment to the Course Allocation instance such that the median budget is $1 + \varepsilon'/2$.

Lemma 8 (Essentially [25]). *Let $n_x > 6\alpha^*$ and suppose that the economy contains the following courses:*

- c_x (the “input course”) ;
- c_{1-x} with capacity $q_{1-x} = 2n_x/3$ (the “output course”);

and the following set of students:

- n_x students interested only in the schedule $\{c_x, c_{1-x}\}$;

and suppose further that at most $n_{1-x} = n_x/6$ other students are interested in course c_{1-x} .

Then in any normalized approximate CEEI with market clearing error at most α^* , at least one of the following must hold:

The gate is ε' -satisfied

$$p_{1-x}^* \in [1 - p_x^*, 1 - p_x^* + \varepsilon']$$

The gadget contributes to income inequality A constant fraction of the n_x students have budgets either less than 1 or greater than $1 + \varepsilon'$.

Proof. Observe that:

- If $p_{1-x}^* > 1 - p_x^* + \varepsilon'$, then none of the n_x students can afford the bundle $\{c_x, c_{1-x}\}$ - except those whose budget is greater than $1 + \varepsilon'$. Other than them, there are at most $n_{1-x} = n_x/6$ students enrolled in the c_{1-x} - much less than the capacity $2n_x/3$. Therefore for the market clearing error to be less than $\alpha^* = n_x/6$, $n_x/3$ of the students must have budget greater than $1 + \varepsilon'$
- On the other hand, if $p_{1-x}^* < 1 - p_x^*$, then all n_x students can afford the bundle $\{c_x, c_{1-x}\}$ - except for those whose budget is less than 1. Therefore in order to satisfy the market clearing requirement, at least $n_x/6$ students must have budget less than 1.

□

Similarly, [25] provide gadgets for all the gates in the definition of the GCIRCUIT problem, and show that they can be concatenated to simulate the computation on the generalized circuit. For each gate of the generalized circuit, [25]’s reduction uses at most $\Theta(1/\log(1/\varepsilon))$ (in particular, a constant) number of gadgets. Furthermore, the number of students that participate in each of those gadgets is approximately the same ($\Theta(\alpha^*)$). Therefore, for every gate which is not ε -satisfied, there are at least $\Theta(\alpha^*)$ students whose budgets are either at most 1 or at least $1 + \varepsilon'$. Thus if the assignment to the Course Allocation problem does not correspond to an (ε, δ) -approximate solution to GCIRCUIT, then a $\delta' = \Theta(\delta/\log(1/\varepsilon))$ -fraction of the students must have budgets at most 1 or at least $1 + \varepsilon'$. Applying Lemma 7 completes the proof of Theorem 4. □